arXiv:1510.00452v4 [cs.LG] 26 Feb 2016

Optimal Binary Classifier Aggregation for General Losses

Akshay Balsubramani
Computer Science and Engineering
University of California, San Diego
abalsubr@ucsd.edu

Yoav Freund
Computer Science and Engineering
University of California, San Diego
yfreund@ucsd.edu

February 29, 2016

Abstract

We address the problem of aggregating an ensemble of binary classifiers in a semi-supervised setting. Recently, this problem was solved optimally by [BF15a] using a game-theoretic approach, but that analysis was specific to the $0-1$ loss. In this paper, we generalize the minimax optimal algorithm of the previous work to a very general, novel class of loss functions, including but not limited to all convex surrogates, while extending its performance and efficiency guarantees.

The result is a family of parameter-free ensemble aggregation algorithms which use labeled and unlabeled data; these are as efficient as linear learning and prediction for convex risk minimization, but work without any relaxations on many non-convex loss functions. The prediction algorithms take a form familiar in decision theory, applying sigmoid functions to a generalized notion of ensemble margin, but without the assumptions typically made in margin-based learning.

1 Introduction

Consider a binary classification problem, in which we attempt to build the best predictor possible for data falling into two classes. At our disposal is an ensemble of individual classifiers which we can use in designing our predictor. The task is to predict with minimum error on a large unlabeled test set, on which we know the predictions of the ensemble classifiers but not the true test labels. This is a prototype supervised learning problem, for which the prototypical supervised solution is to hold out some labeled data to measure the errors of the ensemble classifiers, and then just predict according to the best classifier. But can we use the unlabeled data to better predict using the ensemble classifiers?

This problem is central to semi-supervised learning. It was recently studied by [BF15a], who gave a worst-case-optimal learning and prediction algorithm for it when the evaluation risk, and the constraints, are measured with zero-one classification error. However, the zero-one loss is inappropriate for other common binary classification tasks, such as estimating label probabilities, and handling false positives and false negatives differently. Such goals motivate the use of different losses like log loss and cost-weighted misclassification loss.

In this paper, we generalize the setup of [BF15a] to these loss functions and a large class of others. Like the earlier work, we show that the choice of loss function completely governs an ensemble aggregation algorithm that is minimax optimal in our setting, and is very efficient and scalable to boot.

The algorithm learns a weighting over ensemble classifiers by solving a convex optimization problem. The optimal prediction on each example in the test set turns out to be a sigmoid-like function of a linear combination of the ensemble predictions, using the learned weighting. The minimax structure ensures that this prediction function and the training algorithm are completely data-dependent without parameter choices, relying merely on the structure of the loss function. It also establishes the minimax optimal prediction to have structure reminiscent of a weighted majority vote over the ensemble, and exactly paralleling the prediction function of a generalized linear model ([MN89]).

1.1 Preliminaries

Our setting generalizes that of [BF15a], in which we are given an ensemble $\mathcal{H} = \{h_1, \ldots, h_p\}$ and unlabeled data $x_1, \ldots, x_n$ on which we wish to predict. To start with, the ensemble's predictions on the unlabeled data are denoted by $\mathbf{F}$:

$$\mathbf{F} = \begin{pmatrix} h_1(x_1) & h_1(x_2) & \cdots & h_1(x_n) \\ \vdots & \vdots & \ddots & \vdots \\ h_p(x_1) & h_p(x_2) & \cdots & h_p(x_n) \end{pmatrix} \in [-1,1]^{p \times n} \tag{1}$$

We use vector notation for the rows and columns of $\mathbf{F}$: $\mathbf{h}_i = (h_i(x_1), \cdots, h_i(x_n))^\top$ and $\mathbf{x}_j = (h_1(x_j), \cdots, h_p(x_j))^\top$. The test set has some binary labels $(y_1; \ldots; y_n) \in \{-1, 1\}^n$. As in [BF15a], though, the test labels are allowed to be randomized, represented by values in $[-1, 1]$ instead of just the two values $\{-1, 1\}$. So it is convenient to write the labels on the test data T as $\mathbf{z} = (z_1; \ldots; z_n) \in [-1, 1]^n$. [1] These true test set labels are unknown to the predictor.

Write $[a]_+ = \max(0, a)$ and $[n] = \{1, 2, \ldots, n\}$. All vector inequalities are componentwise. Throughout we use *increasing* to mean "monotonically nondecreasing" to describe a function, and vice versa for *decreasing*.

1.2 Loss Functions

In [BF15a], we suffer loss on each test point according to our expected classification error on it, if our prediction $g_j \in [-1, 1]$ is interpreted as a randomized binary value in the same way as z_j. So if the true label for example j is $y_j = 1$, then the loss of predicting g_j on it is $\ell_+(g_j) = \frac{1}{2}(1 - g_j)$; this is $\{0, 1\}$ when $g_j = \{1, -1\}$ respectively, and a convex combination of the two when $g_j \in (-1, 1)$. Similarly, if $y_j = -1$, then the loss is $\ell_-(g_j) = \frac{1}{2}(1 + g_j)$.

We call $\ell_\pm$ the *partial losses* here, following earlier work (e.g. [RW10]). Since the true label can only be ± 1, the partial losses fully specify the decision-theoretic problem we face, and changing them is tantamount to altering the prediction task.

To guide intuition as to what such partial losses could conceivably look like, observe that they intuitively measure discrepancy to the true label ± 1. As a result it makes sense for e.g. $\ell_+(g)$ to be decreasing, as g increases toward the notional true label $+1$. This suggests that both partial losses $\ell_+(\cdot)$ and $\ell_-(\cdot)$ should be monotonic, which we assume in this paper.

Assumption 1. *Over the interval $(-1, 1)$, $\ell_+(\cdot)$ is decreasing and $\ell_-(\cdot)$ is increasing, and both are twice differentiable.*

We view Assumption 1 as natural, as we have motivated (differentiability is convenient for our proofs, but most of our arguments do not require it; see Section 3 for details). Notably, we do not require convexity or symmetry of the losses. We refer to loss functions whose partial losses satisfying Assumption 1 as "general losses" to contrast them with convex losses or other less broad subclasses – our main learnability result holds for all such losses.

The expected loss we suffer with respect to the *randomized* true labels z_j is a linear combination of the partial losses:

$$\ell(z_j, g_j) := \left(\frac{1 + z_j}{2}\right) \ell_+(g_j) + \left(\frac{1 - z_j}{2}\right) \ell_-(g_j) \tag{2}$$

In decision theory and learning theory, there has been much investigation into the nature of the loss ℓ and its partial losses, particularly on how to estimate the "conditional label probability" z_j using $\ell(z_j, g_j)$. A natural operation to do this is to minimize the loss over g_j; accordingly, a loss ℓ such that $\arg\min_{g \in [-1,1]} \ell(z_j, g) = z_j$ (for all $z_j \in [-1, 1]$) is called a *proper loss* ([BSS05, RW10]), which will be used in later discussions.

[1] For example, a value of $z_i = \frac{1}{2}$ indicates $y_i = +1$ w.p. $\frac{3}{4}$ and -1 w.p. $\frac{1}{4}$.

1.3 Minimax Formulation

The idea of [BF15a] is to formulate the ensemble aggregation problem as a two-player zero-sum game between a predictor and an adversary. In this game, the predictor is the first player, who plays $\mathbf{g} = (g_1; g_2; \ldots; g_n)$, a randomized label $g_j \in [-1, 1]$ for each example $\{x_j\}_{j=1}^n$. The adversary then sets the labels $\mathbf{z} \in [-1, 1]^n$.

The key point is that when any classifier i is known to perform to a certain degree on the test data, its predictions $\mathbf{h}_i$ on the test data are a reasonable guide to $\mathbf{z}$, and correspondingly give us information by constraining $\mathbf{z}$. Each hypothesis in the ensemble contributes to an intersecting set of constraints, which interact in ways that depend on the ensemble predictions $\mathbf{F}$.

Accordingly, for now we assume the predictor has knowledge of a *correlation vector* $\mathbf{b} \in (0, 1]^p$ such that

$$\forall i \in [p], \qquad \frac{1}{n} \sum_{j=1}^{n} h_i(x_j) z_j \geq b_i \tag{3}$$

i.e. $\frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}$. These p inequalities represent upper bounds on individual classifier error rates. We can estimate them from the training set w.h.p. in a statistical learning setting when the training and test data are i.i.d., in a standard way also used by the prototypical supervised empirical risk minimization (ERM) procedure discussed in the introduction ([BF15a]). So in our game-theoretic formulation, the adversary plays under ensemble error constraints defined by $\mathbf{b}$.

The predictor's goal is to *minimize the worst-case expected loss on the test data* (w.r.t. the randomized labeling $\mathbf{z}$), using the loss function $\ell(z, g)$ as defined earlier in Equation (2):

$$\ell(\mathbf{z}, \mathbf{g}) := \frac{1}{n} \sum_{j=1}^{n} \ell(z_j, g_j)$$

This goal can be written as the following optimization problem, a two-player zero-sum game:

$$V := \min_{\mathbf{g} \in [-1,1]^n} \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}}} \ell(\mathbf{z}, \mathbf{g}) \tag{4}$$

$$= \min_{\mathbf{g} \in [-1,1]^n} \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}}} \frac{1}{n} \sum_{j=1}^{n} \left[\left(\frac{1 + z_j}{2} \right) \ell_+(g_j) + \left(\frac{1 - z_j}{2} \right) \ell_-(g_j) \right] \tag{5}$$

In this paper, we solve the learning problem faced by the predictor, finding an optimal strategy $\mathbf{g}^*$ realizing the minimum in (4) for any given "general loss" ℓ. This strategy guarantees good worst-case performance on the unlabeled dataset, with an upper bound of V on the loss. This bound is perfectly tight, by virtue of the argument above – for all $\mathbf{z}_0$ and $\mathbf{g}_0$ obeying the constraints, our definitions give the tight inequalities

$$\min_{\mathbf{g} \in [-1,1]^n} \ell(\mathbf{z}_0, \mathbf{g}) \leq V \leq \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}}} \ell(\mathbf{z}, \mathbf{g}_0) \tag{6}$$

In our formulation of the problem, the constraints on the adversary take a central role. As discussed in previous work with this formulation ([BF15a, BF15b]), these constraints encode the information we have about the true labels. Without them, the adversary would find it optimal to trivially guarantee error (arbitrarily close to) $\frac{1}{2}$ by simply setting all labels uniformly at random ($\mathbf{z} = \mathbf{0}^n$). It is easy to see that adding more information through constraints will never raise the error bound V. [2]

So far, we have given no assumption about the characteristics of $\ell(\mathbf{z}, \mathbf{g})$ other than Assumption 1. Many of our results will require only this, holding for these "general losses."

This brings us to our contributions:

[2] Though it may pose statistical difficulties to do with uniform convergence over the ensemble ([BF15b]).

1. We give the exact minimax $\mathbf{g}^* \in [-1,1]^n$ for general losses (in Section 2.1). The optimal prediction on each example j is a sigmoid function of a fixed linear combination of the ensemble's p predictions on it, so $\mathbf{g}^*$ is a non-convex function of the ensemble predictions on x_j. This is also a constructive proof of a bound on the worst-case loss of *any* predictor constructed from the ensemble, by Equation (6).

2. We derive an efficient algorithm for finding $\mathbf{g}^*$, by solving a p-dimensional convex optimization problem (Section 2.2). We prove this for a broad subclass of losses (the conditions of Lemma 2), including all convex ERM surrogate losses. Extensions to weighted problems and others are in Section 3.

3. The optimal $\mathbf{g}^*$ and an efficient algorithm for it, as above, extended to the case when the constraints can arise from general loss bounds on ensemble classifiers (Section 4), rather than classifier error rates.

2 Results

A few more quantities will be convenient to specify before discussing our main results.

Based on the loss, define the **score function** $\Gamma : [-1,1] \mapsto \mathbb{R}$ is

$$\Gamma(g) := \ell_-(g) - \ell_+(g)$$

(We will also write the vector $\Gamma(\mathbf{g})$ componentwise with $[\Gamma(\mathbf{g})]_j = \Gamma(g_j)$ for convenience, so that $\Gamma(\mathbf{h}_i) \in \mathbb{R}^n$ and $\Gamma(\mathbf{x}_j) \in \mathbb{R}^p$.) Observe that by our assumptions, $\Gamma(g)$ is increasing on its domain, so we can discuss its inverse $\Gamma^{-1}(m)$, which is typically sigmoid-shaped. [3]

With these we will set up the solution to the game (4). The solution depends on the optimum of a convex function, defined here for further use.

Definition 1 (Potential Well)**.** *Define the **potential well***

$$\Psi(m) := \begin{cases} -m + 2\ell_-(-1) & \text{if} \quad m \leq \Gamma(-1) \\ \ell_+(\Gamma^{-1}(m)) + \ell_-(\Gamma^{-1}(m)) & \text{if} \quad m \in (\Gamma(-1), \Gamma(1)) \\ m + 2\ell_+(1) & \text{if} \quad m \geq \Gamma(1) \end{cases} \tag{7}$$

Lemma 2. *The potential well $\Psi(m)$ is continuous and 1-Lipschitz. It is also convex under* any *of the following conditions:*

(A) The partial losses $\ell_\pm(\cdot)$ are convex over $(-1,1)$.
(B) The loss function $\ell(\cdot,\cdot)$ is a proper loss.
(C) $\ell'_-(x)\ell''_+(x) \geq \ell''_-(x)\ell'_+(x)$ for all $x \in (-1,1)$.

(Indeed, the proof shows that the last condition is both sufficient and necessary for convexity of Ψ, under 1.) So the potential wells for different losses are shaped roughly similarly, as seen in Figure 1.

Lemma 2 tells us that the potential well is easy to optimize under any of the given conditions. Note that these conditions encompass convex surrogate losses commonly used in ERM, including all such "margin-based" losses (convex univariate functions of $z_j g_j$). These constitute a large class of losses introduced primarily for their favorable computational properties relative to direct 0-1 loss minimization.

An easily optimized potential well benefits us, because our learning problem basically consists of optimizing it over the unlabeled data, as we will soon make explicit. The function we will actually need to optimize is in terms of the dual parameters, so we call it the slack function.

[3] If Γ does not have a unique inverse, our arguments also work, mutatis mutandis, with the pseudoinverse $\Gamma^{-1}(m) = \inf\{g \in [-1,1] : \Gamma(g) \geq m\}$.

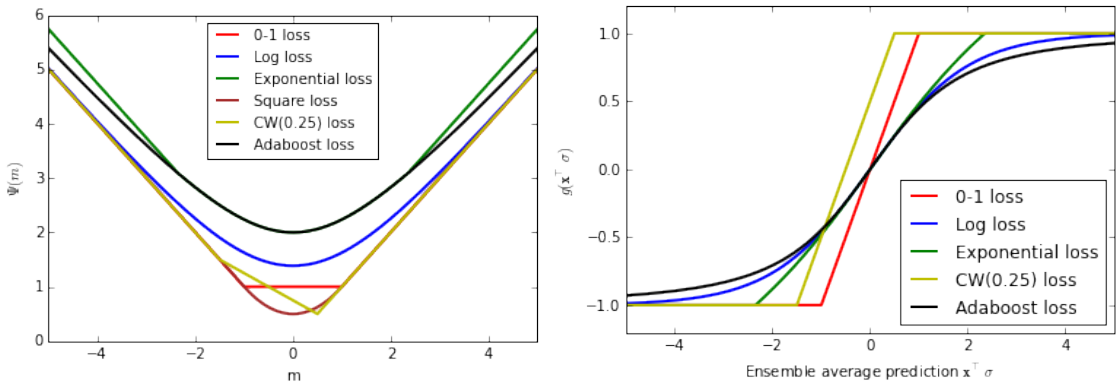

Figure 1: At left are plots of potential wells. At right are optimal prediction functions g, as a function of margin. Both are shown for various losses, as listed in Section 2.3.

Definition 3 (Slack Function)**.** *Let $\sigma \geq 0^p$ be a weight vector over $\mathcal{H}$ (not necessarily a distribution). The vector of **ensemble predictions** is $\mathbf{F}^\top \sigma = (\mathbf{x}_1^\top \sigma, \ldots, \mathbf{x}_n^\top \sigma)$, whose elements' magnitudes are the **margins**. The prediction **slack function** is*

$$\gamma(\sigma, \mathbf{b}) := \gamma(\sigma) := -\mathbf{b}^\top \sigma + \frac{1}{n} \sum_{j=1}^{n} \Psi(\mathbf{x}_j^\top \sigma) \tag{8}$$

*An **optimal weight vector** σ^* is any minimizer of the slack function: $\sigma^* \in \arg\min_{\sigma \geq 0^p} [\gamma(\sigma)]$.*

2.1 Solution of the Game

These are used to describe the minimax equilibrium of the game (4), in our main result.

Theorem 4. *The minimax value of the game* (4) *is*

$$\min_{\mathbf{g} \in [-1,1]^n} \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}}} \ell(\mathbf{z}, \mathbf{g}) = V = \frac{1}{2}\gamma(\sigma^*) = \frac{1}{2} \min_{\sigma \geq 0^p} \left[-\mathbf{b}^\top \sigma + \frac{1}{n} \sum_{j=1}^{n} \Psi(\mathbf{x}_j^\top \sigma) \right]$$

The minimax optimal predictions are defined as follows: for all $j \in [n]$,

$$g_j^* := g_j(\sigma^*) = \begin{cases} -1 & \text{if } \mathbf{x}_j^\top \sigma^* \leq \Gamma(-1) \\ \Gamma^{-1}(\mathbf{x}_j^\top \sigma^*) & \text{if } \mathbf{x}_j^\top \sigma^* \in (\Gamma(-1), \Gamma(1)) \\ 1 & \text{if } \mathbf{x}_j^\top \sigma^* \geq \Gamma(1) \end{cases} \tag{9}$$

g_j^* is always an increasing sigmoid, as shown in Figure 1.

We can also redo the proof of Theorem 4 when $\mathbf{g} \in [-1,1]^n$ is not left as a free variable set in the game, but instead is preset to $\mathbf{g}(\sigma)$ as in (9) for some (possibly suboptimal) weight vector σ.

Observation 5. *For any weight vector $\sigma_0 \geq 0^p$, the worst-case loss after playing $\mathbf{g}(\sigma_0)$ is bounded by*

$$\max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}}} \ell(\mathbf{z}, \mathbf{g}(\sigma_0)) \leq \frac{1}{2}\gamma(\sigma_0)$$

The proof is a simpler version of that of Theorem 4; there is no minimum over $\mathbf{g}$ to deal with, and the minimum over $\sigma \geq 0^p$ in Equation (12) is upper-bounded by using σ_0. This result is an expression of weak duality in our setting, and generalizes Observation 4 of [BF15a].

2.2 The Ensemble Aggregation Algorithm

Theorem 4 defines a prescription for aggregating the given ensemble predictions on the test set. This can be stated in terms of a learning algorithm and a prediction method. Our analysis implies guarantees on the behavior of both.

Learning: *Minimize the slack function $\gamma(\sigma)$, finding the minimizer σ^* that achieves V.*
This is a convex optimization under broad conditions (Lemma 2), and when the test examples are i.i.d. the Ψ term is a sum of n i.i.d. functions. As such it is readily amenable even to standard first-order optimization methods which require only $O(1)$ test examples at once. In practice, learning employs such methods to *approximately* minimize γ, finding some σ_A such that $\gamma(\sigma_A) \leq \gamma(\sigma^*) + \epsilon$ for some small ϵ. Standard convex optimization methods will do this because the slack function is Lipschitz, as Lemma 2 shows (combined with the observation that $\|\mathbf{b}\|_\infty \leq 1$).

Prediction: *Predict $g(\sigma^*)$ on any test example, as indicated in* (9).
This decouples the prediction task on each test example, which is as efficient as p-dimensional linear prediction, requiring $O(p)$ time and memory. After finding an ϵ-approximate minimizer σ_A in the learning step as above, Observation 5 tells us that the prediction $\mathbf{g}(\sigma_A)$ has loss guaranteed to be within $\frac{\epsilon}{2}$ of V.

In particular, note that there is no algorithmic dependence on n in either step in a statistical learning setting, so our transductive formulation is no less tractable than a stochastic optimization setting in which i.i.d. data arrive one at a time.

2.3 Examples of Different Losses

To further illuminate Theorem 4, we detail a few special cases in which ℓ_+, ℓ_- are explicitly defined. These losses may be found throughout the literature; for further information, see [RW10]. The key functions Ψ and g^* are listed for these losses in Table 1, and in most cases plotted in Figure 1. We can see that the nonlinearities used for g^* are all sigmoids, which arises solely from the minimax structure and the box constraints inherent in the classification game (more details in Section 2.4).

- **0-1 Loss**: Here g_j is taken to be a randomized binary prediction; this case was developed in [BF15a], the work we generalize in this paper.
- **Log Loss**
- **Square Loss**
- **Cost-Weighted Misclassification (Quantile) Loss**: This is defined with a parameter $c \in [0, 1]$ representing the relative cost of false positives vs. false negatives, making the Bayes-optimal classifier the c-quantile of the conditional probability distribution ([She05]).
- **Exponential Loss**
- **Logistic Loss**
- **Hellinger Loss**: This is typically given for $p, y \in [0, 1]$ as $\frac{1}{2}\left(\left(\sqrt{p} - \sqrt{y}\right)^2 + \left(\sqrt{1-p} - \sqrt{1-y}\right)^2\right)$. Our formulation is equivalent when the prediction and label are rescaled to $[-1, 1]$.
- **"AdaBoost Loss"**: If the goal of AdaBoost ([SF12]) is interpreted as class probability estimation, the implied loss is proper and given in [BSS05, RW10].
- **Absolute Loss and Hinge Loss**: The absolute loss can be defined by $\ell^{abs}_{\mp}(g_j) = 1 \pm g_j$, and the hinge loss also has $\ell^{abs}_{\mp}(g_j) = 1 \pm g_j$ since the kink in the hinge loss only lies at $g_j = \mp 1$. These partial losses are the same as for 0-1 loss up to scaling, and therefore all our results for Ψ and $\mathbf{g}^*$ are as well. So these losses are not shown in Table 1.

2.4 Proof of Theorem 4

The main hurdle here is the constrained maximization over $\mathbf{z}$. For this we use the following result, a basic application of Lagrange duality (from [BF15a], but proved in Appendix A for completeness).

Lemma 6. *For any* $\mathbf{a} \in \mathbb{R}^n$*,*

$$\max_{\substack{\mathbf{z}\in[-1,1]^n,\\ \frac{1}{n}\mathbf{F}\mathbf{z}\geq\mathbf{b}}} \frac{1}{n}\mathbf{z}^\top\mathbf{a} \;=\; \min_{\sigma\geq 0^p}\left[-\mathbf{b}^\top\sigma + \frac{1}{n}\left\|\mathbf{F}^\top\sigma+\mathbf{a}\right\|_1\right]$$

With this we prove the main result, followed by some comments on the proof.

Proof of Theorem 4. First note that $\ell(\mathbf{z},\mathbf{g})$ is linear in $\mathbf{z}$,

$$V = (5) = \frac{1}{2}\min_{\mathbf{g}\in[-1,1]^n}\ \max_{\substack{\mathbf{z}\in[-1,1]^n,\\ \frac{1}{n}\mathbf{F}\mathbf{z}\geq\mathbf{b}}}\ \frac{1}{n}\sum_{j=1}^n\left[\ell_+(g_j)+\ell_-(g_j)+z_j\left(\ell_+(g_j)-\ell_-(g_j)\right)\right]$$

Here we can rewrite the constrained maximization over $\mathbf{z}$ using Lemma 6:

$$\begin{aligned}\max_{\substack{\mathbf{z}\in[-1,1]^n,\\ \frac{1}{n}\mathbf{F}\mathbf{z}\geq\mathbf{b}}}\ \frac{1}{n}\sum_{i=1}^n z_j\left(\ell_+(g_j)-\ell_-(g_j)\right) &= \max_{\substack{\mathbf{z}\in[-1,1]^n,\\ \frac{1}{n}\mathbf{F}\mathbf{z}\geq\mathbf{b}}}\ -\frac{1}{n}\mathbf{z}^\top[\Gamma(\mathbf{g})] \\ &= \min_{\sigma\geq 0^p}\left[-\mathbf{b}^\top\sigma+\frac{1}{n}\left\|\mathbf{F}^\top\sigma-\Gamma(\mathbf{g})\right\|_1\right] \qquad (10)\end{aligned}$$

Substituting (10) into (5) and simplifying,

$$\begin{aligned}V &= \frac{1}{2}\min_{\mathbf{g}\in[-1,1]^n}\left[\frac{1}{n}\sum_{j=1}^n\left[\ell_+(g_j)+\ell_-(g_j)\right]+\max_{\substack{\mathbf{z}\in[-1,1]^n,\\ \frac{1}{n}\mathbf{F}\mathbf{z}\geq\mathbf{b}}}\ \frac{1}{n}\sum_{j=1}^n z_j\left(\ell_+(g_j)-\ell_-(g_j)\right)\right] \qquad (11)\\
&= \frac{1}{2}\min_{\mathbf{g}\in[-1,1]^n}\left[\frac{1}{n}\sum_{j=1}^n\left[\ell_+(g_j)+\ell_-(g_j)\right]+\min_{\sigma\geq 0^p}\left[-\mathbf{b}^\top\sigma+\frac{1}{n}\left\|\mathbf{F}^\top\sigma-\Gamma(\mathbf{g})\right\|_1\right]\right]\\
&= \frac{1}{2}\min_{\sigma\geq 0^p}\left[-\mathbf{b}^\top\sigma+\min_{\mathbf{g}\in[-1,1]^n}\left[\frac{1}{n}\sum_{j=1}^n\left[\ell_+(g_j)+\ell_-(g_j)\right]+\frac{1}{n}\left\|\mathbf{F}^\top\sigma-\Gamma(\mathbf{g})\right\|_1\right]\right] \qquad (12)\\
&= \frac{1}{2}\min_{\sigma\geq 0^p}\left[-\mathbf{b}^\top\sigma+\frac{1}{n}\sum_{j=1}^n\min_{g_j\in[-1,1]}\left[\ell_+(g_j)+\ell_-(g_j)+\left|\mathbf{x}_j^\top\sigma-\Gamma(g_j)\right|\right]\right] \qquad (13)\end{aligned}$$

The absolute value breaks down into two cases, so the inner minimization's objective can be simplified:

$$\ell_+(g_j)+\ell_-(g_j)+\left|\mathbf{x}_j^\top\sigma-\Gamma(g_j)\right| = \begin{cases}2\ell_+(g_j)+\mathbf{x}_j^\top\sigma & \text{if}\ \ \mathbf{x}_j^\top\sigma\geq\Gamma(g_j)\\ 2\ell_-(g_j)-\mathbf{x}_j^\top\sigma & \text{if}\ \ \mathbf{x}_j^\top\sigma<\Gamma(g_j)\end{cases} \qquad (14)$$

Suppose g_j falls in the first case, so that $\mathbf{x}_j^\top\sigma\geq\Gamma(g_j)$. From Assumption 1, $2\ell_+(g_j)+\mathbf{x}_j^\top\sigma$ is decreasing in g_j, so it is minimized for the greatest $g_j^*\leq 1$ s.t. $\Gamma(g_j^*)\leq\mathbf{x}_j^\top\sigma$. Since $\Gamma(\cdot)$ is increasing, exactly one of two subcases holds:

1. g_j^* is such that $\Gamma(g_j^*)=\mathbf{x}_j^\top\sigma$, in which case the minimand (14) is $\ell_+(g_j^*)+\ell_-(g_j^*)$
2. $g_j^*=1$ so that $\Gamma(g_j^*)=\Gamma(1)<\mathbf{x}_j^\top\sigma$, in which case the minimand (14) is $2\ell_+(1)+\mathbf{x}_j^\top\sigma$

A precisely analogous argument holds if g_j falls in the second case, where $\mathbf{x}_j^\top \sigma < \Gamma(g_j)$. Putting the cases together, we have shown the form of the summand Ψ, piecewise over its domain, so (13) is equal to $\frac{1}{2}\min_{\sigma \geq 0^p}[\gamma(\sigma)]$.

We have proved the dependence of g_j^* on $\mathbf{x}_j^\top \sigma^*$, where σ^* is the minimizer of the outer minimization of (13). This completes the proof. □

There are two notable ways in which the result of Theorem 4 is particularly convenient and general. First, the fact that $\ell(z, g)$ can be non-convex in g, yet solvable as a convex problem, is a major departure from previous work. Second, the solution has a convenient dependence on n like in [BF15a], simply averaging a function over the unlabeled data, which is not only mathematically convenient but also makes stochastic $O(1)$-space optimization practical. This is surprisingly powerful, because the original minimax problem is jointly over the entire dataset, avoiding further independence or decoupling assumptions.

Both these favorable properties stem from the structure of the binary classification problem, as we can describe by examining the optimization problem we constructed in Equation (13) of this proof. In it, the constraints which do not explicitly appear with Lagrange parameters are all box, or L_∞ norm, constraints. These decouple over the n test examples, which is a critical reason we are able to obtain results for such general non-convex problems. It allows us to reduce the problem to the one-dimensional optimization at the heart of Equation (13), which we can solve with ad hoc methods. When we do so, the g_i are "clipped" sigmoid functions because of the bounding effect of the $[-1, 1]$ constraint and loss function structure.

We introduce Lagrange parameters σ only for the p remaining constraints in the problem, which do not so decouple, applying globally over the n test examples. However, these constraints only depend on n as an average over examples (which how they arise in dual form in Equation (22) of the proof), and the loss function itself is also an average (Equation (11)). Putting these together shows how the favorable dependence on n comes about here, making an efficient solution possible to a potentially flagrantly non-convex problem. The technique of optimizing only "halfway into" the dual allows us to readily manipulate the minimax problem exactly without using an approximation like weak duality, despite the lack of convexity in $\mathbf{g}$.

A final bit of structure in the binary classification problem – the linearity of $\ell(z, g)$ in z – is used in Section 4 to make a further sweeping generalization of Theorem 4.

3 Related Work and Extensions

A number of further comments are in order.

The predecessor to this work ([BF15a]) addresses a problem, 0-1 loss minimization, that is well known to be strongly NP-hard when solved directly. Formulating it in the transductive setting, in which the data distribution is known, is crucial. It gives the dual problem an independently interesting interpretation, so the learning problem is on the always-convex Lagrange dual function and is therefore tractable.

This work generalizes that idea, as the possibly non-convex partial losses are minimized transductively via a straightforward convex optimization. A similar formal technique, including the crucial use of L_∞-norm constraints and averaging over $n(\gg p)$ to decompose optimization efficiently over many examples, is used implicitly, for a different purpose, in the "drifting" repeated game analysis of boosting ([SF12], Sec. 13.4.1). Existing boosting work ([SF12]) is loosely related to our approach in being a transductive game invoked to analyze ensemble aggregation, but it does not consider unlabeled data and draws its power instead from being a repeated game.

Our transductive formulation involves no surrogates or relaxations of the loss, which we view as a significant advantage over previous work – it allows us to bypass the consistency and agnostic-learning discussions ([Zha04, BJM06]) common to ERM methods that use convex risk minimization. Convergence analyses of such methods make heavy use of convexity of the losses and are generally done presupposing a linear weighting on $h \in \mathcal{H}$ ([TDS15]), whereas in our work such structure emerges directly from Lagrange duality and involves no convexity to derive the worst-case-optimal predictions. Prior work does express the idea that is our explicit conclusion – the learning problem is completely determined by the choice of loss function.

The conditions in 1 are notable for their generality. Differentiability is convenient, but by no means necessary. We only choose to use it because first-order conditions are a convenient way to establish convexity of the potential well in Lemma 2. It is never used elsewhere, including in the minimax argument of Theorem 4.

Also, we assert that the monotonicity condition on $\ell_\pm$ is natural for a loss function. If it were otherwise, for instance if there were $g_1 < g_2 < g_3$ s.t. $\ell_-(g_1) = \ell_-(g_3) \neq \ell_-(g_2)$, then this would indicate that the loss function is not simply a function of the disparity between labels; in this case perhaps a different loss function is warranted. In technical terms, the minimax manipulations we use to prove Theorem 4 are structured to be valid *even if $\ell_\pm$ are non-monotonic*; but in this case, g_j^* could turn out to be multi-valued and hence not a true function of the example's margin. So we consider there to be significant evidence that our assumption on the losses is necessary, but do not prove it here. We note with interest that the same family of general losses was defined by [She05] in the ERM setting (dubbed "F-losses" there) – in which condition (C) of Lemma 2 also has significance ([She05], Prop. 2) – but has seemingly not been revisited thereafter.

Indeed, we emphasize that our result on the minimax equilibrium (Theorem 4) holds for all such general losses – the slack function may not be convex unless the further conditions of Lemma 2 are met, but the interpretation of the optimum in terms of margins and sigmoid functions remains. All this structure emerges from the inherent decision-theoretic structure of the problem. We have already discussed this in Section 2.4, but another example: it is easy to see that the function $g(\mathbf{x}_j^\top \sigma)$ is always increasing in $\mathbf{x}_j^\top \sigma$ for general losses, because the score function Γ is increasing. This monotonicity typically needs to be assumed in generalized linear models ([MN89]) and related settings, which use such objects for prediction. The score function can be thought of as analogous to the link function used by GLMs (log loss corresponds to the scaled and shifted logit link, as seen from Table 1), with its inverse being used for prediction. Further fleshing out the links between our minimax analysis and GLM learning would be interesting future work.

All our algorithms in this manuscript can be used in full generality with "specialist" hypotheses in the ensemble that only predict on some subset of the test examples. This is done by merely calculating the loss bounds only over these examples, and changing $\mathbf{F}$ and $\mathbf{b}$ accordingly; see [BF15b], which also shows the experimental efficacy of this approach.

3.1 Weighted Test Sets, Covariate Shift, and Label Noise

Though our results here deal with binary classification of a uniformly-weighted test set, note that our formulation deals with a weighted test set with only a modification to the slack function:

Theorem 7. *For any vector $\mathbf{r} \geq \mathbf{0}^n$,*

$$\min_{\mathbf{g}\in[-1,1]^n} \max_{\substack{\mathbf{z}\in[-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z}\geq\mathbf{b}}} \frac{1}{n}\sum_{j=1}^n r_j \ell(z_j, g_j) = \frac{1}{2}\min_{\sigma\geq 0^p}\left[-\mathbf{b}^\top\sigma + \frac{1}{n}\sum_{j=1}^n r_j \Psi\left(\frac{\mathbf{x}_j^\top\sigma}{r_j}\right)\right]$$

With $\sigma_{\mathbf{r}}^*$ defined as the minimizer to the above, the optimal predictions $\mathbf{g}^* = \mathbf{g}(\sigma_{\mathbf{r}}^*)$, as in Theorem 4.

Such weighted classification can be parlayed into algorithms for general supervised learning problems via learning reductions ([BLZ08]). Allowing weights on the test set for the evaluation is tantamount to accounting for known covariate shift in our setting; it would be interesting, though outside our scope, to investigate scenarios with more uncertainty.

In addition, varying the weights used in Theorem 7 can be interpreted as changing box constraints on $\mathbf{z}$; defining $\tilde{\mathbf{z}} = \mathbf{r} \circ \mathbf{z}$, we have

$$\min_{\mathbf{g}\in[-1,1]^n} \max_{\substack{\mathbf{z}\in[-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z}\geq\mathbf{b}}} \frac{1}{n}\sum_{j=1}^n r_j \ell(z_j, g_j) = \min_{\mathbf{g}\in[-1,1]^n} \max_{\substack{-\mathbf{r}\leq\tilde{\mathbf{z}}\leq\mathbf{r}, \\ \frac{1}{n}\tilde{\mathbf{F}}\tilde{\mathbf{z}}\geq\mathbf{b}}} \ell(\tilde{\mathbf{z}}, \mathbf{g})$$

where $\tilde{\mathbf{F}}$ is a suitably redefined version of $\mathbf{F}$ (s.t. $\tilde{\mathbf{x}}_j = \frac{1}{r_j}\mathbf{x}_j$). The right-hand side here is formally equivalent to the original problem except for the box constraint on the adversary, which is now nonuniform. This was

done for the 0-1 loss in ([BF15a], Prop. 5-6), where it was interpreted as constraining the adversary to act under a level of known label noise when $\|\mathbf{r}\|_\infty \leq 1$. It is clear from the above that when $\|\mathbf{r}\|_\infty \leq 1$,

$$\min_{\mathbf{g}\in[-1,1]^n} \max_{\substack{-\mathbf{r}\leq\mathbf{z}\leq\mathbf{r},\\ \frac{1}{n}\mathbf{F}\mathbf{z}\geq\mathbf{b}}} \ell(\mathbf{z},\mathbf{g}) = \frac{1}{2}\min_{\sigma\geq 0^p}\left[-\mathbf{b}^\top\sigma + \frac{1}{n}\sum_{j=1}^n r_j\Psi\left(\mathbf{x}_j^\top\sigma\right)\right] \leq \frac{1}{2}\min_{\sigma\geq 0^p}\left[\gamma(\sigma)\right] = V$$

i.e. knowing the noise level always helps in a minimax sense by further constraining $\mathbf{z}$, as was seen for the 0-1 loss in [BF15a].

3.2 Uniform Convergence Bounds for $\mathbf{b}$

Given $\mathbf{b}$ as a lower bound on ensemble classifier losses, the slack function can be efficiently approximately optimized, translating into a worst-case prediction loss bound, as we have seen in Section 2.2. But there is typically error in estimating $\mathbf{b}$, often quantified by uniform convergence (L_∞) bounds on $\mathbf{b}$. If one would like to incorporate the two-sided constraint on $\mathbf{b}$, the solution to our game involves the dual (L_1) norm of the dual vector σ.

Theorem 8. *We have*

$$\min_{\mathbf{g}\in[-1,1]^n} \max_{\substack{\mathbf{z}\in[-1,1]^n,\\ \left\|\frac{1}{n}\mathbf{F}\mathbf{z}-\mathbf{b}\right\|_\infty\leq\epsilon}} \ell(\mathbf{z},\mathbf{g}) \;=\; \min_{\sigma\in\mathbb{R}^p}\left[-\mathbf{b}^\top\sigma + \frac{1}{n}\sum_{j=1}^n \Psi(\mathbf{x}_j^\top\sigma) + \epsilon\left\|\sigma\right\|_1\right]$$

Let σ_s^* be the minimizer of the right-hand side above. Then the optimal $\mathbf{g}^* = \mathbf{g}(\sigma_s^*)$, the same function of the optimal weighting as in (9)*.*

(This is proved exactly like Theorem 4, but using Lemma 13 instead of Lemma 6 in that proof.) So when the ensemble losses are uniformly bounded, we are now searching over all vectors σ (not just nonnegative ones) in an L_1-regularized version of the original optimization problem in Theorem 4. To our knowledge, this particular analysis has been addressed in prior work only for the special case of the entropy objective function on the probability simplex – [DPS04] discusses that special case further.

4 Constraints on General Losses

In the previous sections, we allow the evaluation function of the game to be a general loss, but assume that the constraints (our information about the ensemble) are in terms of zero-one loss. Here we relax that assumption, allowing each classifier h_i to constrain the test labels $\mathbf{z}$, not with the zero-one loss of h_i's predictions, but rather with some other general loss.

This is possible because when the true labels are binary, all the losses we consider are linear in $\mathbf{z}$, as seen in (5): $\ell(\mathbf{z},\mathbf{g}) = \frac{1}{n}\sum_{j=1}^n \frac{1}{2}\left[\ell_+(g_j) + \ell_-(g_j)\right] - \frac{1}{2n}\mathbf{z}^\top[\Gamma(\mathbf{g})]$. Accordingly, recall that $\mathbf{h}_i \in [-1,1]^n$ is the vector of test predictions of hypothesis h_i. Suppose we have an upper bound on the generalization loss of h_i, i.e. $\ell(\mathbf{z},\mathbf{h}_i) \leq \epsilon_i^\ell$. If we define

$$b_i^\ell := \frac{1}{n}\sum_{j=1}^n \left[\ell_+(h_i(x_j)) + \ell_-(h_i(x_j))\right] - 2\epsilon_i^\ell \tag{15}$$

then we can use the definition of $\ell(\mathbf{z},\mathbf{g})$ to write

$$\ell(\mathbf{z},\mathbf{h}_i) \leq \epsilon_i^\ell \qquad \Longleftrightarrow \qquad \frac{1}{n}\mathbf{z}^\top[\Gamma(\mathbf{h}_i)] \geq b_i^\ell \tag{16}$$

Now (16) is a linear constraint on $\mathbf{z}$, just like each of the error constraints earlier considered in (3). We can derive an aggregation algorithm with constraints like (16), using essentially the same analysis as employed in Section 2 to solve the game (4).

In summary, any classifier can be incorporated into our framework for aggregation if we have a generalization loss bound on it, where the loss can be any of the losses we have considered. This allows an enormous variety of constraint sets, as each classifier considered can have constraints corresponding to any number of loss bounds on it, including 0 (it can be omitted from the constraint set, providing no information about $\mathbf{z}$), or > 1 (it can conceivably have multiple loss bounds using different losses). For instance, h_1 can yield a constraint corresponding to a zero-one loss bound, h_2 can yield one constraint corresponding to a square loss bound and another corresponding to a zero-one loss bound, and so on.

4.1 Matching Objective and Constraint Losses

Despite this generality, we can glean some intuition about the aggregation method for general losses. To do so in the rest of this section, we only consider the case when each classifier contributes exactly one constraint to the problem, and the losses used for these constraints are all the same as each other and as the loss ℓ used in the objective function. In other words, the minimax prediction problem we consider is

$$V^\ell := \min_{\mathbf{g}\in[-1,1]^n} \max_{\substack{\mathbf{z}\in[-1,1]^n,\\ \forall i\in[p]:\ \ell(\mathbf{z},\mathbf{h}_i)\le\epsilon_i^\ell}} \ell(\mathbf{z},\mathbf{g}) = \min_{\mathbf{g}\in[-1,1]^n} \max_{\substack{\mathbf{z}\in[-1,1]^n,\\ \forall i\in[p]:\ \frac{1}{n}\mathbf{z}^\top[\Gamma(\mathbf{h}_i)]\ge b_i^\ell}} \ell(\mathbf{z},\mathbf{g}) \tag{17}$$

The matrix $\mathbf{F}$ and the slack function from (1) are therefore redefined:

$$F_{ij}^\ell := \Gamma(h_i(x_j)) = \ell_-(h_i(x_j)) - \ell_+(h_i(x_j)) \tag{18}$$

$$\gamma^\ell(\sigma, \mathbf{b}^\ell) := \gamma^\ell(\sigma) := -[\mathbf{b}^\ell]^\top\sigma + \frac{1}{n}\sum_{j=1}^n \Psi\left([\Gamma(\mathbf{x}_j)]^\top\sigma\right) \tag{19}$$

where $\mathbf{b}^\ell = (b_1^\ell, \ldots, b_p^\ell)^\top$ and the vectors $\mathbf{x}_j$ are now from the new unlabeled data matrix F_{ij}^ℓ. The game (17) is clearly of the same form as the earlier formulation (4). Therefore, its solution has the same structure as in Theorem 4, proved using that theorem's proof:

Theorem 9. *The minimax value of the game* (17) *is* $V := \frac{1}{2}\gamma^\ell(\sigma^{\ell*}) := \min_{\sigma\ge\mathbf{0}^p} \frac{1}{2}\gamma^\ell(\sigma)$. *The minimax optimal predictions are defined as follows: for all* $j \in [n]$,

$$g_j^* := g_j(\sigma^*) = \begin{cases} -1 & \text{if } [\Gamma(\mathbf{x}_j)]^\top\sigma^{\ell*} \le \Gamma(-1) \\ \Gamma^{-1}\left([\Gamma(\mathbf{x}_j)]^\top\sigma^{\ell*}\right) & \text{if } [\Gamma(\mathbf{x}_j)]^\top\sigma^{\ell*} \in (\Gamma(-1), \Gamma(1)) \\ 1 & \text{if } [\Gamma(\mathbf{x}_j)]^\top\sigma^{\ell*} \ge \Gamma(1) \end{cases} \tag{20}$$

This provides a concise characterization of how to solve the semi-supervised binary classification game for general losses. Though on the face of it Theorem 9 is a much stronger result than even Theorem 4, we cannot neglect statistical issues. The loss bounds ϵ_i^ℓ on each classifier are estimated using a uniform convergence bound over the ensemble with loss ℓ, but the data now considered are not the ensemble predictions, but the predictions passed through the score function Γ. This can be impractical for losses like log loss, for which Γ is unbounded, and therefore uniform convergence to estimate b_i^ℓ in (16) is much looser than for 0-1 loss.

But such issues are outside our scope here, and our constrained minimax results hold in any case. They may be useful to obtain semi-supervised learnability results for different losses from tighter statistical characterizations, which we consider an interesting open problem.

4.2 Beating the Best Classifier and the Best Weighted Majority

As in the predecessor work [BF15a, BF15b], our guarantee here is the strongest possible worst-case guarantee, because our method is the minimax algorithm. This type of statement lacks precedent in this area, so we conclude our discussions by highlighting a few simple corollaries when the loss is general.

In minimizing the slack function over the dual parameters σ, we perform at least as well as the weighting $\sigma^i \geq \mathbf{0}^p$ that puts weight 1 on h_i and 0 on the remaining classifiers $h_{i' \neq i}$. In other words, our predictor always has the option of simply choosing the best single classifier i^* and guaranteeing its loss bound $\epsilon^\ell_{i^*}$. Consequently, our predictor's loss is always at most that of any single classifier, proving the following observation.

Proposition 10. *$V^\ell \leq \epsilon^\ell_i$ for any classifier $i \in [p]$ and any loss ℓ.*

We provide a short alternative proof of this in Appendix A using the definitions of this section to better illuminate how they fit together, even though the proposition is evident from the fact that we are minimizing over $\{\sigma : \sigma \geq 0^p\} \ni \sigma^i$.

Our minimax guarantee means that given the ensemble loss constraints $\mathbf{b}^\ell$, our algorithm automatically admits superior worst-case loss bounds to any predictor, including *any weighted majority vote* as well.

Proposition 11. *V^ℓ has a better expected test loss guarantee than any weighted majority vote.*

The precise meaning of Proposition 11 should be stressed. It is saying that any weighted majority vote could have a $\mathbf{z}$, consistent with the information in the ensemble constraints, that leads to more errors than the error guarantee for $\mathbf{g}^*$. It does not imply that the actual performance of $\mathbf{g}^*$ will beat every weighted majority, because the constraints could still be unhelpful and the bounds loose.

Recent empirical results of [BF15b] suggest that this and the statistical issues we have discussed in estimating $\mathbf{b}$ are major considerations in the empirical performance of such algorithms, but can be dealt with to realize significant practical benefits from unlabeled data.

References

[BF15a] Akshay Balsubramani and Yoav Freund. Optimally combining classifiers using unlabeled data. In *Conference on Learning Theory*, 2015.

[BF15b] Akshay Balsubramani and Yoav Freund. Scalable semi-supervised classifier aggregation. In *Advances in Neural Information Processing Systems*, 2015.

[BJM06] Peter L Bartlett, Michael I Jordan, and Jon D McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.

[BLZ08] Alina Beygelzimer, John Langford, and Bianca Zadrozny. Machine learning techniques?reductions between prediction quality metrics. In *Performance Modeling and Engineering*, pages 3–28. Springer, 2008.

[BSS05] Andreas Buja, Werner Stuetzle, and Yi Shen. Loss functions for binary class probability estimation and classification: Structure and applications. 2005.

[DPS04] Miroslav Dudik, Steven J Phillips, and Robert E Schapire. Performance guarantees for regularized maximum entropy density estimation. In *Learning Theory*, pages 472–486. Springer, 2004.

[MN89] Peter McCullagh and John A Nelder. *Generalized linear models*, volume 37. CRC press, 1989.

[RW10] Mark D Reid and Robert C Williamson. Composite binary losses. *The Journal of Machine Learning Research*, 11:2387–2422, 2010.

[SF12] Robert E. Schapire and Yoav Freund. *Boosting: Foundations and Algorithms*. The MIT Press, 2012.

[She05] Yi Shen. *Loss functions for binary classification and class probability estimation*. PhD thesis, University of Pennsylvania, 2005.

[TDS15] Matus Telgarsky, Miroslav Dudik, and Robert Schapire. Convex risk minimization and conditional probability estimation. In *Conference on Learning Theory*, 2015.

[Zha04] Tong Zhang. Statistical behavior and consistency of classification methods based on convex risk minimization. *Annals of Statistics*, pages 56–85, 2004.

A Supporting Results and Proofs

Proof of Lemma 6. We have

$$\max_{\substack{\mathbf{z}\in[-1,1]^n,\\ \mathbf{F}\mathbf{z}\geq n\mathbf{b}}} \frac{1}{n}\mathbf{z}^\top\mathbf{a} = \frac{1}{n}\max_{\mathbf{z}\in[-1,1]^n}\min_{\sigma\geq 0^p}\left[\mathbf{z}^\top\mathbf{a}+\sigma^\top(\mathbf{F}\mathbf{z}-n\mathbf{b})\right] \tag{21}$$

$$\overset{(a)}{=} \frac{1}{n}\min_{\sigma\geq 0^p}\max_{\mathbf{z}\in[-1,1]^n}\left[\mathbf{z}^\top(\mathbf{a}+\mathbf{F}^\top\sigma)-n\mathbf{b}^\top\sigma\right] \tag{22}$$

$$= \frac{1}{n}\min_{\sigma\geq 0^p}\left[\left\|\mathbf{a}+\mathbf{F}^\top\sigma\right\|_1 - n\mathbf{b}^\top\sigma\right] = \min_{\sigma\geq 0^p}\left[-\mathbf{b}^\top\sigma+\frac{1}{n}\left\|\mathbf{F}^\top\sigma+\mathbf{a}\right\|_1\right] \tag{23}$$

where (a) is by the minimax theorem ([SF12], p.144). □

Lemma 12. *The function $\ell_+(\Gamma^{-1}(m))+\ell_-(\Gamma^{-1}(m))$ is convex for $m\in(\Gamma(-1),\Gamma(1))$ under any of the conditions of Lemma 2.*

Proof of Lemma 12. Define $F(m)=\ell_+(\Gamma^{-1}(m))+\ell_-(\Gamma^{-1}(m))$. By basic properties of the derivative,

$$\frac{d\left[\Gamma^{-1}\right]}{dm} = \frac{1}{\Gamma'(\Gamma^{-1}(m))} = \frac{1}{\ell'_-(\Gamma^{-1}(m))-\ell'_+(\Gamma^{-1}(m))} \geq 0 \tag{24}$$

where the last inequality follows by Assumption 1. Therefore, by the chain rule and (24),

$$F'(m) = \frac{\ell'_-(\Gamma^{-1}(m))+\ell'_+(\Gamma^{-1}(m))}{\ell'_-(\Gamma^{-1}(m))-\ell'_+(\Gamma^{-1}(m))} \tag{25}$$

From this, we calculate $F''(m)$, writing $\ell'_\pm(\Gamma^{-1}(m))$ and $\ell''_\pm(\Gamma^{-1}(m))$ as simply $\ell'_\pm$ and $\ell''_\pm$ for clarity:

$$F''(m) = \underbrace{\frac{\frac{d\left[\Gamma^{-1}\right]}{dm}}{\left(\ell'_-(\Gamma^{-1}(m))-\ell'_+(\Gamma^{-1}(m))\right)^2}}_{(a)}\left[\left(\ell'_--\ell'_+\right)\left(\ell''_-+\ell''_+\right)-\left(\ell'_-+\ell'_+\right)\left(\ell''_--\ell''_+\right)\right]$$

From (24), observe that the term $(a)=\left(\ell'_-(\Gamma^{-1}(m))-\ell'_+(\Gamma^{-1}(m))\right)^{-3}\geq 0$. Therefore, it suffices to show that the other term is ≥ 0. But this is equal to

$$\left(\ell'_--\ell'_+\right)\left(\ell''_-+\ell''_+\right)-\left(\ell'_-+\ell'_+\right)\left(\ell''_--\ell''_+\right) = 2(\ell'_-\ell''_+-\ell''_-\ell'_+) \tag{26}$$

This proves that condition (C) of Lemma 2 is sufficient for convexity of F (and indeed necessary also, under 1 on the partial losses).

We now address the other conditions of Lemma 2. (A) implies (C), because by Assumption 1, $\ell'_-, \ell''_+, \ell''_-$ are ≥ 0 and $\ell'_+ \leq 0$, so (26) is ≥ 0 as desired.

Finally we prove that (B) implies (C). If ℓ is proper, then it is well known (e.g. Thm. 1 of [RW10], and [BSS05]) that for all $x \in (-1, 1)$,

$$\frac{\ell'_-(x)}{1+x} = -\frac{\ell'_+(x)}{1-x}$$

(This is a simple and direct consequence of stationary conditions from the properness definition.)

Define the function $w(x) = \frac{\ell'_-(x)}{1+x} = -\frac{\ell'_+(x)}{1-x}$; we drop the argument and write it and its derivative as w and w' for clarity. By direct computation,

$$\begin{aligned}\ell'_-\ell''_+ - \ell''_-\ell'_+ &= [(1+x)w\,(w + (x-1)w')] - [(w + (1+x)w')(x-1)w] \\ &= \left[(1+x)w^2 + (x^2-1)ww'\right] - \left[(x-1)w^2 + (x^2-1)ww'\right] = 2w^2 \geq 0\end{aligned}$$

so (C) is true as desired. □

Proof of Lemma 2. Continuity follows by checking $\Psi(m)$ at $m = \pm 1$.

For Lipschitzness, note that for $m \in (\Gamma(-1), \Gamma(1))$, by (25),

$$\Psi'(m) = \frac{\ell'_-(\Gamma^{-1}(m)) + \ell'_+(\Gamma^{-1}(m))}{\ell'_-(\Gamma^{-1}(m)) - \ell'_+(\Gamma^{-1}(m))} \tag{27}$$

$$= -1 + \frac{2\ell'_-(\Gamma^{-1}(m))}{\ell'_-(\Gamma^{-1}(m)) - \ell'_+(\Gamma^{-1}(m))} \tag{28}$$

$$= 1 - \frac{2(-\ell'_+(\Gamma^{-1}(m)))}{\ell'_-(\Gamma^{-1}(m)) - \ell'_+(\Gamma^{-1}(m))} \tag{29}$$

Using Assumption 1 on the partial losses, equations (28) and (29) respectively make clear that $\Psi'(m) \geq -1$ and $\Psi'(m) \leq 1$ on this interval. Since $\Psi'(m)$ is -1 for $m < \Gamma(-1)$ and 1 for $m > \Gamma(1)$, it is 1-Lipschitz.

As for convexity, since Ψ is linear outside the interval $(\Gamma(-1), \Gamma(1))$, it suffices to show that $\Psi(m)$ is convex inside this interval, which is shown in Lemma 12. □

Lemma 13. *For any* $\mathbf{a} \in \mathbb{R}^n$*,*

$$\max_{\substack{\mathbf{z}\in[-1,1]^n, \\ \left\|\frac{1}{n}\mathbf{F}\mathbf{z}-\mathbf{b}\right\|_\infty \leq \epsilon}} \frac{1}{n}\mathbf{z}^\top \mathbf{a} \;=\; \min_{\sigma\in\mathbb{R}^p} \left[-\mathbf{b}^\top\sigma + \frac{1}{n}\left\|\mathbf{F}^\top\sigma + \mathbf{a}\right\|_1 + \epsilon\left\|\sigma\right\|_1\right]$$

Proof.

$$\begin{aligned}\max_{\substack{\mathbf{z}\in[-1,1]^n, \\ \left\|\frac{1}{n}\mathbf{F}\mathbf{z}-\mathbf{b}\right\|_\infty \leq \epsilon}} \frac{1}{n}\mathbf{z}^\top\mathbf{a} &= \max_{\substack{\mathbf{z}\in[-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z}-\mathbf{b}\leq\epsilon\mathbf{1}^n, \\ -\frac{1}{n}\mathbf{F}\mathbf{z}+\mathbf{b}\leq\epsilon\mathbf{1}^n}} \frac{1}{n}\mathbf{z}^\top\mathbf{a} \\ &= \frac{1}{n}\max_{\mathbf{z}\in[-1,1]^n}\min_{\lambda,\xi\geq 0^p}\left[\mathbf{z}^\top\mathbf{a} + \lambda^\top(-\mathbf{F}\mathbf{z} + n\mathbf{b} + n\epsilon\mathbf{1}^n) + \xi^\top(\mathbf{F}\mathbf{z} - n\mathbf{b} + n\epsilon\mathbf{1}^n)\right] \\ &= \frac{1}{n}\min_{\lambda,\xi\geq 0^p}\max_{\mathbf{z}\in[-1,1]^n}\left[\mathbf{z}^\top(\mathbf{a} + \mathbf{F}^\top(\xi-\lambda)) + \lambda^\top(n\mathbf{b} + n\epsilon\mathbf{1}^n) + \xi^\top(-n\mathbf{b} + n\epsilon\mathbf{1}^n)\right] \\ &= \frac{1}{n}\min_{\lambda,\xi\geq 0^p}\left[\left\|\mathbf{a} + \mathbf{F}^\top(\xi-\lambda)\right\|_1 - n\mathbf{b}^\top(\xi-\lambda) + n\epsilon\mathbf{1}^{n\top}(\xi+\lambda)\right]\end{aligned}$$

Suppose for some $j \in [n]$ that $\xi_j > 0$ and $\lambda_j > 0$. Then subtracting $\min(\xi_j, \lambda_j)$ from both does not affect the value $[\xi - \lambda]_j$, but always decreases $[\xi + \lambda]_j$, and therefore always decreases the objective function.

Therefore, we can w.l.o.g. assume that $\forall j \in [n] : \min(\xi_j, \lambda_j) = 0$. Defining $\sigma_j = \xi_j - \lambda_j$ for all j (so that $\xi_j = [\sigma_j]_+$ and $\lambda_j = [\sigma_j]_-$), the last equality above becomes

$$\frac{1}{n} \min_{\lambda,\xi \geq 0^p} \left[\left\| \mathbf{a} + \mathbf{F}^\top (\xi - \lambda) \right\|_1 - n\mathbf{b}^\top (\xi - \lambda) + n\epsilon \mathbf{1}^{n\top} (\xi + \lambda) \right] = \frac{1}{n} \min_{\sigma \in \mathbb{R}^p} \left[\left\| \mathbf{a} + \mathbf{F}^\top \sigma \right\|_1 - n\mathbf{b}^\top \sigma + n\epsilon \left\| \sigma \right\|_1 \right]$$

□

Proof of Theorem 7. The proof is quite similar to that of Theorem 4, but generalizes it. First note that we have

$$\begin{aligned} \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}}} \frac{1}{n} \sum_{i=1}^n r_j z_j \left(\ell_+(g_j) - \ell_-(g_j) \right) &= \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}}} -\frac{1}{n} \mathbf{z}^\top [\mathbf{r} \circ \Gamma(\mathbf{g})] \\ &= \min_{\sigma \geq 0^p} \left[-\mathbf{b}^\top \sigma + \frac{1}{n} \left\| \mathbf{F}^\top \sigma - (\mathbf{r} \circ \Gamma(\mathbf{g})) \right\|_1 \right] \end{aligned} \tag{30}$$

where the last equality uses Lemma 6.

Therefore, using (30) on the left-hand side of what we wish to prove,

$$\begin{aligned} V &= \frac{1}{2} \min_{\mathbf{g} \in [-1,1]^n} \left[\frac{1}{n} \sum_{j=1}^n r_j \left[\ell_+(g_j) + \ell_-(g_j) \right] + \max_{\substack{\mathbf{z} \in [-1,1]^n, \\ \frac{1}{n}\mathbf{F}\mathbf{z} \geq \mathbf{b}}} \frac{1}{n} \sum_{i=1}^n r_j z_j \left(\ell_+(g_j) - \ell_-(g_j) \right) \right] \\ &= \frac{1}{2} \min_{\mathbf{g} \in [-1,1]^n} \left[\frac{1}{n} \sum_{j=1}^n r_j \left[\ell_+(g_j) + \ell_-(g_j) \right] + \min_{\sigma \geq 0^p} \left[-\mathbf{b}^\top \sigma + \frac{1}{n} \sum_{j=1}^n \left| \mathbf{x}_j^\top \sigma - r_j \Gamma(g_j) \right| \right] \right] \\ &= \frac{1}{2} \min_{\sigma \geq 0^p} \left[-\mathbf{b}^\top \sigma + \frac{1}{n} \sum_{j=1}^n \min_{g_j \in [-1,1]} \left(r_j \left[\ell_+(g_j) + \ell_-(g_j) \right] + \left| \mathbf{x}_j^\top \sigma - r_j \Gamma(g_j) \right| \right) \right] \end{aligned} \tag{31}$$

As in the proof of Theorem 4, the inner minimization's objective can be simplified:

$$r_j(\ell_+(g_j) + \ell_-(g_j)) + \left| \mathbf{x}_j^\top \sigma - r_j \Gamma(g_j) \right| = \begin{cases} 2r_j \ell_+(g_j) + \mathbf{x}_j^\top \sigma & \text{if } \mathbf{x}_j^\top \sigma \geq r_j \Gamma(g_j) \\ 2r_j \ell_-(g_j) - \mathbf{x}_j^\top \sigma & \text{if } \mathbf{x}_j^\top \sigma < r_j \Gamma(g_j) \end{cases} \tag{32}$$

Suppose g_j falls in the first case, so that $\mathbf{x}_j^\top \sigma \geq r_j \Gamma(g_j)$. From Assumption 1, $2r_j \ell_+(g_j) + \mathbf{x}_j^\top \sigma$ is decreasing in g_j, so it is minimized for the greatest $g_j^* \leq 1$ s.t. $\Gamma(g_j^*) \leq \frac{\mathbf{x}_j^\top \sigma}{r_j}$. Since $\Gamma(\cdot)$ is increasing, exactly one of two subcases holds:

a) g_j^* is such that $\Gamma(g_j^*) = \frac{\mathbf{x}_j^\top \sigma}{r_j}$, in which case the minimand (32) is $r_j(\ell_+(g_j^*) + \ell_-(g_j^*))$

b) $g_j^* = 1$ so that $\Gamma(g_j^*) = \Gamma(1) < \frac{\mathbf{x}_j^\top \sigma}{r_j}$, in which case the minimand (32) is $2r_j \ell_+(1) + \mathbf{x}_j^\top \sigma$

A precisely analogous argument holds if g_j falls in the second case, where $\mathbf{x}_j^\top \sigma < \Gamma(g_j)$. So as before, we have proved the dependence of g_j^* on $\mathbf{x}_j^\top \sigma^*$, where σ^* is the minimizer of the outer minimization of (31). This completes the proof. □

Proof of Prop. 10. Consider the weighting σ^i as above. Then

$$V^\ell = \frac{1}{2} \min_{\sigma \geq \mathbf{0}^p} \gamma^\ell(\sigma) \leq \frac{1}{2} \gamma^\ell(\sigma^i) = -b_i^\ell + \frac{1}{n} \sum_{j=1}^n \Psi(\Gamma(h_i(x_j))) \tag{33}$$

Since $h_i(x_j) \in [-1,1] \quad \forall j$, we have $\Gamma(h_i(x_j)) \in [\Gamma(-1), \Gamma(1)]$. Using the definitions of Ψ and $\mathbf{b}^\ell$ (in Equation (15)), (33) can therefore be rewritten as

$$\begin{aligned} 2V^\ell &\leq -b_i^\ell + \frac{1}{n}\sum_{j=1}^{n}\left[\ell_+(\Gamma^{-1}(\Gamma(h_i(x_j)))) + \ell_-(\Gamma^{-1}(\Gamma(h_i(x_j))))\right] \\ &= -\frac{1}{n}\sum_{j=1}^{n}[\ell_+(h_i(x_j)) + \ell_-(h_i(x_j))] + 2\epsilon_i^\ell + \frac{1}{n}\sum_{j=1}^{n}[\ell_+(h_i(x_j)) + \ell_-(h_i(x_j))] = 2\epsilon_i^\ell \end{aligned}$$

□

Loss	Partial Losses	$\Gamma(g)$	$\Psi(m)$	$g_i(\sigma)$
0-1	$\ell_-(g) = \frac{1}{2}(1+g)$ $\ell_+(g) = \frac{1}{2}(1-g)$	g	$\max(1, \lvert m \rvert)$	$\mathrm{clip}(\mathbf{x}_i^\top \sigma)$
Log	$\ell_-(g) = \ln\left(\frac{2}{1-g}\right)$ $\ell_+(g) = \ln\left(\frac{2}{1+g}\right)$	$\ln\left(\frac{1+g}{1-g}\right)$	$\ln(1+e^m) + \ln(1+e^{-m})$	$\frac{1-e^{-\mathbf{x}_i^\top \sigma}}{1+e^{-\mathbf{x}_i^\top \sigma}}$
Square	$\ell_-(g) = \left(\frac{1+g}{2}\right)^2$ $\ell_+(g) = \left(\frac{1-g}{2}\right)^2$	g	$\begin{cases} -m & m \le -1 \\ \frac{1}{2}(m^2+1) & m \in (-1,1) \\ m & m \ge 1 \end{cases}$	$\mathrm{clip}(\mathbf{x}_i^\top \sigma)$
CW (param. c)	$\ell_-(g) = c(1+g)$ $\ell_+(g) = (1-c)(1-g)$	$g + 2c - 1$	$\begin{cases} -m & m \le 2c-2 \\ (2c-1)m + 4c(1-c) & m \in (2c-2, 2c) \\ m & m \ge 2c \end{cases}$	$\mathrm{clip}(\mathbf{x}_i^\top \sigma + 1 - 2c)$
Exponential	$\ell_-(g) = e^g$ $\ell_+(g) = e^{-g}$	$e^g - e^{-g}$	$\begin{cases} -m + 2/e & m \le -e + \frac{1}{e} \\ \sqrt{4+m^2} & m \in \left(-e+\frac{1}{e}, e-\frac{1}{e}\right) \\ m + 2/e & m \ge e - \frac{1}{e} \end{cases}$	$\mathrm{clip}\left(\ln\left(\frac{1}{2}\mathbf{x}_i^\top \sigma + \sqrt{1 + \frac{1}{4}(\mathbf{x}_i^\top \sigma)^2}\right)\right)$
Logistic	$\ell_-(g) = \ln(1+e^g)$ $\ell_+(g) = \ln\left(1+e^{-g}\right)$	g	$\begin{cases} -m + 2\ln(1+1/e) & m \le -1 \\ \ln(1+e^m) + \ln(1+e^{-m}) & m \in (-1,1) \\ m + 2\ln(1+1/e) & m \ge 1 \end{cases}$	$\mathrm{clip}(\mathbf{x}_i^\top \sigma)$
Hellinger	$\ell_-(g) = 1 - \sqrt{\frac{1-g}{2}}$ $\ell_+(g) = 1 - \sqrt{\frac{1+g}{2}}$	$\sqrt{\frac{1+g}{2}} - \sqrt{\frac{1-g}{2}}$	$\begin{cases} -m & m \le -1 \\ 2 - \left(\frac{1-m\sqrt{2-m^2}}{2}\right)^{1/2} - \left(\frac{1+m\sqrt{2-m^2}}{2}\right)^{1/2} & m \in (-1,1) \\ m & m \ge 1 \end{cases}$	$\begin{cases} (\mathbf{x}_i^\top \sigma)\sqrt{2 - (\mathbf{x}_i^\top \sigma)^2} & \lvert \mathbf{x}_i^\top \sigma \rvert \le 1 \\ \mathrm{sgn}(\mathbf{x}_i^\top \sigma) & \lvert \mathbf{x}_i^\top \sigma \rvert > 1 \end{cases}$
"AdaBoost"	$\ell_-(g) = \sqrt{\frac{1+g}{1-g}}$ $\ell_+(g) = \sqrt{\frac{1-g}{1+g}}$	$\frac{2g}{\sqrt{1-g^2}}$	$\left(\frac{\sqrt{m^2+4}+m}{\sqrt{m^2+4}-m}\right)^{1/2} + \left(\frac{\sqrt{m^2+4}-m}{\sqrt{m^2+4}+m}\right)^{1/2}$	$\frac{\mathbf{x}_i^\top \sigma}{\sqrt{(\mathbf{x}_i^\top \sigma)^2+4}}$

Table 1: Some losses for transductive classifier aggregation, as in Sec. 2.3. For convenience, we write $\mathrm{clip}(x) = \min(1, \max(-1, x))$.